

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/46		A2	(11) International Publication Number: WO 99/44133 (43) International Publication Date: 2 September 1999 (02.09.99)
(21) International Application Number: PCT/US99/04065 (22) International Filing Date: 25 February 1999 (25.02.99) (30) Priority Data: 60/076,048 26 February 1998 (26.02.98) US 09/045,652 20 March 1998 (20.03.98) US (71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, MS UPAL01-521, Palo Alto, CA 94303 (US). (72) Inventors: SCHEIFLER, Robert; 96 North Street, Somerville, MA 02144 (US). WOLLRATH, Ann, M.; 9 Northwoods Road, Groton, MA 01450 (US). JONES, Peter, C.; 85 Bacon Street, Winchester, MA 01890 (US). (74) Agents: GARRETT, Arthur, S.; Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P., 1300 I Street, N.W., Washington, DC 20005-3315 (US) et al.			(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>Without international search report and to be republished upon receipt of that report.</i>
(54) Title: METHOD AND SYSTEM FOR DETERMINISTIC HASHES TO IDENTIFY REMOTE METHODS			
(57) Abstract <p>A method and system is provided to uniquely identify a remote method to invoke on a server using a hash value computed from the method signature sent from the client to the server with the call request. When a client wishes to invoke a remote method located on a server, the client sends a hash value identifying the remote method to the server in the "remote method invocation" (RMI) call. In one implementation, this hash value is created by applying a hash function to the method string name and the parameter type list and possibly the return type. When the server receives the RMI call, the server identifies which method is being called using the received hash value. The server maintains a mapping of hash values to their associated remote methods located on the server and references the correct method using the hash value. Additionally, in one implementation, the server creates the mapping table dynamically when a remote object is created. The server identifies the methods implemented by the object and creates hash values for each method. These hash values are stored in a mapping table which is used to reference the remote methods.</p>			



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		



METHOD AND SYSTEM FOR DETERMINISTIC
HASHES TO IDENTIFY REMOTE METHODS

Related Applications

The following identified U.S. patent applications are relied upon and are incorporated by reference in this application.

5 Provisional U.S. Patent Application No. 60/076,048, entitled "Distributed Computing System," filed on February 26, 1998.

 U.S. Patent Application No. 09/044,923, entitled "Method and System for Leasing Storage," bearing attorney docket no. 06502.0011-01000, and filed on the same date herewith.

10 U.S. Patent Application No. 09/044,838, entitled "Method, Apparatus, and Product for Leasing of Delegation Certificates in a Distributed System," bearing attorney docket no. 06502.0011-02000, and filed on the same date herewith.

 U.S. Patent Application No. 09/044,834, entitled "Method, Apparatus and Product for Leasing of Group Membership in a Distributed System," bearing attorney docket no. 06502.0011-03000, and filed on the same date herewith.

15 U.S. Patent Application No. 09/044,916, entitled "Leasing for Failure Detection," bearing attorney docket no. 06502.0011-04000, and filed on the same date herewith.

 U.S. Patent Application No. 09/044,933, entitled "Method for Transporting Behavior in Event Based System," bearing attorney docket no. 06502.0054-00000, and filed on the same date herewith.

20 U.S. Patent Application No. 09/044,919, entitled "Deferred Reconstruction of Objects and Remote Loading for Event Notification in a Distributed System," bearing attorney docket no. 06502.0062-01000, and filed on the same date herewith.

 U.S. Patent Application No. 09/044,938, entitled "Methods and Apparatus for Remote Method Invocation," bearing attorney docket no. 06502.0102-00000, and filed on the same date herewith.

25 U.S. Patent Application No. 09/044,790, entitled "Method and Apparatus for Determining Status of Remote Objects in a Distributed System," bearing attorney docket no. 06502.0104-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,930, entitled "Downloadable Smart Proxies for Performing Processing Associated with a Remote Procedure Call in a Distributed System," bearing attorney docket no. 06502.0105-00000, and filed on the same date herewith.

5 U.S. Patent Application No. 09/044,917, entitled "Suspension and Continuation of Remote Methods," bearing attorney docket no. 06502.0106-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,835, entitled "Method and System for Multi-Entry and Multi-Template Matching in a Database," bearing attorney docket no. 06502.0107-00000, and filed on the same date herewith.

10 U.S. Patent Application No. 09/044,839, entitled "Method and System for In-Place Modifications in a Database," bearing attorney docket no. 06502.0108, and filed on the same date herewith.

15 U.S. Patent Application No. 09/044,945, entitled "Method and System for Typesafe Attribute Matching in a Database," bearing attorney docket no. 06502.0109-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,931, entitled "Dynamic Lookup Service in a Distributed System," bearing attorney docket no. 06502.0110-00000, and filed on the same date herewith.

20 U.S. Patent Application No. 09/044,939, entitled "Apparatus and Method for Providing Downloadable Code for Use in Communicating with a Device in a Distributed System," bearing attorney docket no. 06502.0112-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,826, entitled "Method and System for Facilitating Access to a Lookup Service," bearing attorney docket no. 06502.0113-00000, and filed on the same date herewith.

25 U.S. Patent Application No. 09/044,932, entitled "Apparatus and Method for Dynamically Verifying Information in a Distributed System," bearing attorney docket no. 06502.0114-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/030,840, entitled "Method and Apparatus for Dynamic Distributed Computing Over a Network," and filed on February 26, 1998.

U.S. Patent Application No. 09/044,936, entitled "An Interactive Design Tool for Persistent Shared Memory Spaces," bearing attorney docket no. 06502.0116-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,934, entitled "Polymorphic Token-Based Control,"
5 bearing attorney docket no. 06502.0117-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,915, entitled "Stack-Based Access Control," bearing attorney docket no. 06502.0118-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,944, entitled "Stack-Based Security Requirements,"
10 bearing attorney docket no. 06502.0119-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,837, entitled "Per-Method Designation of Security Requirements," bearing attorney docket no. 06502.0120-00000, and filed on the same date herewith.

Background

A. Field of the invention

15 This invention relates to data processing systems, and more particularly to remote method invocations on remote servers. Even more specifically, this invention relates to a method and system for identifying remote methods on a server machine using hash values.

B. Related Art

20 Distributed systems typically comprise multiple machines, such as computers and related peripheral devices, connected in a network, for example, a Local Area Networks (LAN), Wide Area Network (WAN), or the Internet. Distributed systems generally require that computational entities (*e.g.*, applications, programs, applets, etc.) running in different address spaces, potentially on different machines, be able to communicate.

25 For a basic communication mechanism, distributed object oriented systems utilize a Remote Procedure Call (RPC) mechanism referred to as Remote Method Invocation (RMI). RMI facilitates application-level communication between "objects" residing in different address spaces. In object oriented systems, a "class" provides a template for the creation of "objects" (which represent items or instances manipulated by the system) having characteristics of that
30 class. The term template denotes that the objects (*i.e.*, data items) in each class, share certain

characteristics or attributes determined by the class such as its methods. Objects are typically created dynamically during system operation. Methods associated with a class are generally invoked (*i.e.*, caused to operate) on the objects of the same class.

RMI is the action of invoking a method of a remote object. In response to the invocation of a method of a remote object using RMI, a lower level communications process causes the invoked method to be executed on the remote object.

The Java™ runtime system, which is designed to implement applications written in the Java™ object oriented programming language, supports a specific Java™ RMI Application Program Interface (API). This API is explained in, for example, a document entitled "Remote Method Invocation Specification," Sun Microsystems, Inc. (1997), which is available via universal resource locator (URL) <http://www.javasoft.com/products/jdk/1.1/docs/guide/rmi/spec/rmiTOC.doc.html>, and is incorporated herein by reference. The Java™ language is described in many texts, including one that is entitled "The Java Language Specification" by James Gosling, Bill Joy, and Guy Steele, Addison-Wesley, 1996. Java and all Java-based trademarks are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Java RMI assumes a homogeneous environment of the specialized Java runtime system, and therefore Java RMI takes advantage of a specialized object model for the Java language whenever possible. In the Java™ distributed object model, a remote object is one that has methods that can be invoked from another runtime system, potentially on a different machine. An object of this type is described by one or more remote interfaces code written in the Java language that specify the methods of the remote object.

The Java runtime system keeps track of all remote objects referenced by computational entities executing through a local virtual machine (VM). The Java™ VM (JVM) is an abstract computing machine of the runtime system that receives instructions from programs in the form of bytecodes and that interprets these bytecodes by dynamically converting them into a form for execution, such as object code, and executing them. The JVM is described in detail in a text entitled "The Java Virtual Machine Specification", by Tim Lindholm and Frank Yellin, Addison Wesley, 1996.

In Java RMI, a client, while processing a program, can remotely initiate processing by a server computer of methods in connection with certain "parameter" information provided by the client. After the server has processed the procedure, it will provide results of its processing to the client, which the client may thereafter use in its processing operations. Typically in such RMI calls, the client will make use of a local "stub" which, when called, transfers the request to the server which implements the particular method, obtains the results and provides them back to the client.

Conventionally, when a client calls a method on a remote object containing a list of methods, the method is identified by a string name or a sequence number identifying the selected method. However, identifying a method by its string name can create false identification of a remote method because the remote object may have more than one method with the same string name. Such methods are said to be "overloaded." Although methods may have the same string name, overloaded methods with duplicate same string names typically take different parameter types. For instance, suppose a remote object, using the Java programming language, has the following two methods:

```
public interface Directory {  
    PhoneNumber lookupPhone(String name);  
    PhoneNumber lookupPhone(Person person)  
}
```

If a client seeks to invoke one of these two methods, the string name "lookupPhone" alone does not enable the remote object to determine the correct method to be invoked because more than one method with that name exist.

Another conventional approach for identifying a remote method is to put the methods in alphabetical order and number them. Suppose a remote object implements the following two methods:

```
public interface Directory {  
    PhoneNumber lookupPhone(String name);  
    void storePhone(String name, PhoneNumber phone);  
}
```

The numbering of the methods may be represented as follows:

1. lookupPhone
2. storePhone

When a client wants to invoke a method, it simply sends the number corresponding to the method in the method invocation instruction. If, however, new methods are added to the remote object such that it appears as follows:

```
public interface Directory {  
    PhoneNumber lookupPhone(String name);  
    void storePhone(String name, PhoneNumber phone);  
    Address lookupAddress(String name);  
    void storeAddress(String name, Address addr);  
}
```

the new numbering of the methods would be:

1. lookupAddress
2. lookupPhone
3. storeAddress
4. storePhone

Hence, the numbers corresponding to each method have changed. Thus, existing clients that continue to use old stubs using the old numbering would invoke the wrong methods.

Accordingly, it is desirable to provide a system that uniquely identifies the methods of remote objects for RMI.

Summary

The present invention satisfies this and other desires by providing a method and system for identifying the methods of remote objects using hash values.

A method in a data processing system for invoking a remote methods comprises the steps of providing a hash value uniquely identifying a remote method, sending the hash value in response to an instruction to invoke the remote method, and invoking the remote method based on the hash value. The method further includes the step of locating the remote method in a mapping table using the hash value.

Apparatus and systems are also provided for carrying out methods consistent with the present invention.

The advantages accruing to the present invention are numerous. For example, methods and systems consistent with the present invention identify unique remote methods for invocation, thus avoiding false identification of incorrect remote methods. Furthermore, this identification can be performed even if two or more methods have the same string name, or the methods use a changing numbering system.

Although a long string such as the method name combined with a parameter type list could be used to more precisely identify a remote method, such an identifier would be cumbersome. The use of hash values further creates greater efficiency by eliminating the need for long strings to more precisely identify remote methods. Additionally, it allows the server to perform more efficiently because the server can more efficiently manipulate and compute the integer numbers than the strings.

It is, therefore, desirable to provide a method and apparatus to uniquely identify remote methods using hash values.

Brief Description of the Drawings

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate an embodiment of the invention and, together with the description, serve to explain the advantages and principles of the invention. In the drawings,

Figure 1 illustrates a network in which systems consistent with the present invention may be implemented;

Figure 2 is block diagram of the system architecture for a computer system with which the present invention may be implemented;

Figure 3 is a block diagram illustrating an RMI call using a hash value between a client computer and a server computer consistent with the present invention;

Figure 4 is a block diagram of a hash value mapping table consistent with the present invention;

Figure 5 is a flowchart illustrating the steps used to identify a unique remote method consistent with the present invention; and

Figure 6 is a flowchart illustrating the steps used by a server machine to create a hash value mapping table consistent with the present invention.

Detailed Description

Overview

Methods and systems consistent with the present invention identify a method of a remote object using a hash value. When a client wishes to invoke a method of a remote object located on a server, the client sends the hash value identifying the particular remote method to the server over the RMI connection. In one implementation, this hash value is created by applying a hash function to the method string name and the parameter type list. Known hash functions with low collision rates can be used for this purpose.

When the server receives the method invocation, the server identifies the called method using the received hash value. The server maintains a mapping of hash values to their associated remote methods located on the server and references the correct method using the hash value.

The server creates the mapping table dynamically when a remote object is created. Upon the creation of a remote object, hash values are determined for each method implemented by the remote object. The server then adds these hash values and pointers to their corresponding methods to the mapping table. When adding the hash value and method pointer, the server checks the mapping table to verify that the pairing is unique, *i.e.*, the server checks for a hash value collision. This process allows remote methods to be identified uniquely and allows the server to continually add methods over time, as the remote class evolves, without notifying all clients with old stubs of the new methods. Additionally, it allows clients using old stubs to correctly identify remote methods on the server. Even further, the use of hashes avoids the need for long strings to identify remote methods.

The Distributed System

Methods and systems consistent with the present invention operate in a distributed system ("the exemplary distributed system") with various components, including both hardware and software. The exemplary distributed system (1) allows users of the system to share services and resources over a network of many devices; (2) provides programmers with tools and programming patterns that allow development of robust, secured distributed systems; and (3)

simplifies the task of administering the distributed system. To accomplish these goals, the exemplary distributed system utilizes the Java™ programming environment to allow both code and data to be moved from device to device in a seamless manner. Accordingly, the exemplary distributed system is layered on top of the Java programming environment and exploits the characteristics of this environment, including the security offered by it and the strong typing provided by it. The Java programming environment is more clearly described in Jaworski, Java 1.1 Developer's Guide, Sams.net (1997), which is incorporated herein by reference.

In the exemplary distributed system, different computers and devices are federated into what appears to the user to be a single system. By appearing as a single system, the exemplary distributed system provides the simplicity of access and the power of sharing that can be provided by a single system without giving up the flexibility and personalized response of a personal computer or workstation. The exemplary distributed system may contain thousands of devices operated by users who are geographically disperse, but who agree on basic notions of trust, administration, and policy. Within the exemplary distributed system are various logical groupings of services provided by one or more devices, and each such logical grouping is known as a Djinn. A "service" refers to a resource, data, or functionality that can be accessed by a user, program, device, or another service and that can be computational, storage related, communication related, or related to providing access to another user. Examples of services provided as part of a Djinn include devices, such as printers, displays, and disks; software, such as applications or utilities; information, such as databases and files; and users of the system.

Both users and devices may join a Djinn. When joining a Djinn, the user or device adds zero or more services to the Djinn and may access, subject to security constraints, any one of the services it contains. Thus, devices and users federate into a Djinn to share access to its services. The services of the Djinn appear programmatically as objects of the Java programming environment, which may include other objects, software components written in different programming languages, or hardware devices. A service has an interface defining the operations that can be requested of that service, and the type of the service determines the interfaces that make up that service.

Fig. 1 depicts the exemplary distributed system 100 containing a computer 102, a computer 104, and a device 106 interconnected by a network 108. The device 106 may be any

of a number of devices, such as a printer, fax machine, storage device, computer, or other devices. The network 108 may be a local area network, wide area network, or the Internet. Although only two computers and one device are depicted as comprising the exemplary distributed system 100, one skilled in the art will appreciate that the exemplary distributed system 100 may include additional computers or devices.

Fig. 2 depicts the computer 102 in greater detail to show a number of the software components of the exemplary distributed system 100. One skilled in the art will appreciate that computer 104 or device 106 may be similarly configured. Computer 102 includes a memory 202, a secondary storage device 204, a central processing unit (CPU) 206, an input device 208, and a video display 210. The memory 202 includes a lookup service 212, a discovery server 214, and a Java™ runtime system 216. The Java runtime system 216 includes the Java™ remote method invocation system (RMI) 218 and a Java™ virtual machine 220. The secondary storage device 204 includes a Java™ space 222.

As mentioned above, the exemplary distributed system 100 is based on the Java programming environment and thus makes use of the Java runtime system 216. The Java runtime system 216 includes the Java™ API, allowing programs running on top of the Java runtime system to access, in a platform-independent manner, various system functions, including windowing capabilities and networking capabilities of the host operating system. Since the Java API provides a single common API across all operating systems to which the Java runtime system 216 is ported, the programs running on top of a Java runtime system run in a platform-independent manner, regardless of the operating system or hardware configuration of the host platform. The Java runtime system 216 is provided as part of the Java™ software development kit available from Sun Microsystems of Mountain View, CA.

The Java virtual machine 220 also facilitates platform independence. The Java virtual machine 220 acts like an abstract computing machine, receiving instructions from programs in the form of byte codes and interpreting these byte codes by dynamically converting them into a form for execution, such as object code, and executing them. RMI 218 facilitates remote method invocation by allowing objects executing on one computer or device to invoke methods of an object on another computer or device. Both RMI and the Java virtual machine are also provided as part of the Java software development kit.

The lookup service 212 defines the services that are available for a particular Djinn. That is, there may be more than one Djinn and, consequently, more than one lookup service within the exemplary distributed system 100. The lookup service 212 contains one object for each service within the Djinn, and each object contains various methods that facilitate access to the corresponding service. The lookup service 212 and its access are described in greater detail in co-pending U.S. Patent Application No. _____, entitled "Method and System for Facilitating Access to a Lookup Service," which has previously been incorporated by reference.

The discovery server 214 detects when a new device is added to the exemplary distributed system 100, during a process known as boot and join or discovery, and when such a new device is detected, the discovery server passes a reference to the lookup service 212 to the new device, so that the new device may register its services with the lookup service and become a member of the Djinn. After registration, the new device becomes a member of the Djinn, and as a result, it may access all the services contained in the lookup service 212. The process of boot and join is described in greater detail in co-pending U.S. Patent Application No. _____, entitled "Apparatus and Method for providing Downloadable Code for Use in Communicating with a Device in a Distributed System," which has previously been incorporated by reference.

The Java space 222 is an object repository used by programs within the exemplary distributed system 100 to store objects. Programs use the Java space 222 to store objects persistently as well as to make them accessible to other devices within the exemplary distributed system. Java spaces are described in greater detail in co-pending U.S. Patent Application No. 08/971,529, entitled "Database System Employing Polymorphic Entry and Entry Matching," assigned to a common assignee, filed on November 17, 1997, which is incorporated herein by reference. One skilled in the art will appreciate that the exemplary distributed system 100 may contain many lookup services, discovery servers, and Java spaces.

Although systems and methods consistent with the present invention are described as operating in the exemplary distributed system and the Java programming environment, one skilled in the art will appreciate that the present invention can be practiced in other systems and other programming environments. Additionally, although aspects of the present invention are described as being stored in memory, one skilled in the art will appreciate that these aspects can

also be stored on or read from other types of computer-readable media, such as secondary storage devices, like hard disks, floppy disks, or CD-ROM; a carrier wave from the Internet; or other forms of RAM or ROM. Sun, Sun Microsystems, the SunLogo, Java, and Java-based trademarks are trademarks or registered trademarks of Sun Microsystems Inc. in the United States and other countries.

Identifying Remote Methods Using Hashes

Figure 3 is a block diagram illustrating an RMI call using a hash value consistent with the present invention. It also shows two computers, client 302 and server 312, which may correspond to computers 102 and 104 shown in distributed system 100. The invocation of a method on a remote object is implemented using Java RMI, although other RMI mechanisms may be used. When client 302 wishes to access a method implemented by a remote object 314 on a server 312, client 302 uses a stub 304 referencing remote object 314. Stub 304 is typically downloaded from server 312 but can also be local to the client 302 or downloaded from somewhere else in network 100, including another server. The manner in which the client obtains a stub is described in greater detail in copending U.S. Patent Application No. 08/636,706, entitled, "System and Method For Facilitating Dynamic Loading of 'Stub' Information to Enable a Program Operating in One Address Space to Invoke Processing of a Remote Method or Procedure in Another Address Space", herein incorporated by reference. Additionally, a "stubless" implementation may be employed in a manner consistent with U.S. Patent Application No. _____, entitled "Methods and Apparatus for Remote Method Invocation", bearing attorney docket no. 06502.0102-00000, which was previously incorporated by reference.

Stub 304, which references remote object 314, has a local method 306 for each remote method, such as remote method 316, implemented by remote object 314. This local method 306 is implemented by the client to invoke the corresponding method 316. It performs functions, such as initiating the communication link between the client and remote method 316 and sending the hash value identifying the method. It should be noted, however, that remote object 314 may have more than one method, although only one method 316 is shown in Figure 3. Similarly, stub

304 may have more than one local method to implement remote methods, but only one is shown in Figure 3 for simplicity.

In one implementation consistent with the present invention, local method 306 is created during the compiling of stub 304, which is created by server 312. When a user supplies a remote object 316 (in the form of Java source code) as a Java class, a Java compiler (not shown) on server 312 compiles the Java class, thus creating a binary class file. This binary class file is compiled by a stub compiler (not shown) on server to create a stub class. Clients use instances of this stub class (*i.e.*, a stub) to invoke methods of the remote object 316.

In this implementation, local method 306 is compiled into stub 304 during the process of compiling the stub. The stub compiler compiles hash value 308 into the local method. As a result, local method 306 has a hash value that identifies the corresponding method in the remote object referenced by the stub. For example, suppose a server has a remote method:

```
int insurancePremium (String state, int age)
```

Then the corresponding stub may have a local method implemented as follows:

```
int insurancePremium (String state, int age) {  
    Stream out = startNewCall ( );  
    sendLong (out, 4056878021019060934...);  
    sendString (out, state);  
    sendInt (out, age);  
    Stream in = finishCall (out);  
    String result = readString (in);  
    finishResults (in);  
    return result;  
}
```

where the long integer of the sendLong method call is a hash value uniquely identifying a remote method.

In one implementation, hash value 308 is a hash value resulting from applying a standard hash function to the combination of the method name and parameter type list 318 of the remote method 316, as follows:

Hash (Method Name, Parameter Type List)

This hash function returns a hash value that may be an integer. Both the method name and parameter type list are used to avoid collisions otherwise caused by using only the method name alone. In another implementation, the hash function may be applied to the method name, parameter type and return value type. In other implementations, however, the hash function may be applied to the method name alone where collisions are less likely.

In another implementation consistent with the present invention, the hash function applied to the method name and parameter type list is the hash function used by the "Secure Hash Algorithm 1" (SHA-1) secure hash standard. This standard is described in detail in the Federal Information Processing Standard Publication 180-1, "Secure Hash Standard", and can also be found at <http://csrc.nist.gov/fips> on the Internet. The SHA-1 protocol is a hash function that produces a 160 bit hash value. In yet another implementation consistent with the present invention, the hash value used is only the first 64 bits of the SHA-1 hash value. In this implementation, the hash value 308 is represented by these 64 bits, a shorten version of the full SHA-1 hash value.

Figure 3 also shows RMI call 310, which is used when client 302 sends a message to invoke a remote method on a remote server such as server 312. RMI call 310 further includes hash value 308. Upon receipt of RMI call 310, server 312 then uses the hash value 308 to reference mapping table 320 and identify a selected remote method.

Figure 4 further depicts details of mapping table 320 on server 312 consistent with the present invention. Generally, mapping table 320 represents the mapping of hash values to individual remote methods of a remote object 314 on server 312. As such, mapping table 320 includes sets of pairings 402 of a hash value 404 and a pointer to a remote method 306. This pointer to a method is a "handle" that identifies a method in such a way as to allow it to be programmatically invoked through the handle. For example, in the Java programming language, this would be an instance of `java.lang.reflect.Method`. In C++, it would be a function pointer (*i.e.*, the actual machine address of the code). As a result, each hash value 304 references a remote method 306.

(1) Identifying Remote Methods

Figure 5 illustrates the steps used in a method consistent with the present invention for identifying a unique remote method on a server by using hash values. First, client 302 makes an RMI call 310 to server 312 to remotely invoke a remote method 316 on server 312. In this RMI call 310, client 302 sends a hash value identifying remote method 316 to be invoked (step 500). In RMI call 310, the client 302 may also pass any parameter arguments to be used by the remote method 316 to invoke the method.

Next, server 312 receives hash value 308 included in RMI call 310 (step 502). Server 312 then accesses mapping table 320 for the server class of remote object 314 to identify which remote method is to be invoked (step 504). Upon accessing mapping table 320, server 312 uses hash value 308 sent in RMI call 310 to identify the remote method to be invoked in the mapping table.

At this point, server 312 invokes method 316 using the received parameter argument values in RMI call 310 (step 506). Finally, server 312 returns the result of the method invocation to client 302 (step 508).

For an example using these steps of a method consistent with the present invention, suppose a remote object implemented the following exemplary set of methods:

```
public interface Directory {
    PhoneNumber lookupPhone(String name);
    PhoneNumber lookupPhone(Person person);
    void storePhone(String name, PhoneNumber phone);
    void storePhone(Person person, PhoneNumber phone);
    Address lookupAddress(String name);
    Address lookupAddress(Person person);
    void storeAddress(String name, Address addr);
    void storeAddress(Person person, Address addr);
}
```

Because this list of remote methods includes methods with duplicate string names, accessing the list by method name may result in invocation of the wrong method. If, for instance, a client wished to invoke the first lookupPhone method listed in the example, the client would send an RMI call including the hash of the method name and parameter type list:

Hash (lookupPhone, String)

This process ensures that the second method, lookupPhone with the parameter Person, would not be invoked. In addition to this hash, the client also sends the argument for the parameter String (*i.e.*, "John" to lookup the phone number for a person with the string name John.)

(2) Building the Mapping Table

Figure 6 depicts the steps used in methods consistent with the present invention by server 312 for dynamically building the mapping table 320 at run time. Generally, when a remote object is created, the Java runtime system on server 312 adds the hash values for each method of the remote object to the mapping table 320. As a result, server 312 has a mapping table for each remote class, since typically all remote objects of the same class have the same remote methods.

First, in methods consistent with the present invention, an object on server 312 is created as remote object, such as object 314 (step 600). Upon this creation, the Java runtime system on server 312 locates all remote methods 316 supported by object 314 (step 602). The Java runtime system calculates the hash value for each remote method 316 of the remote object 314. In one implementation, it obtains the method name and parameter type list 318 (step 604) and computes the hash of the method name and parameter type list (step 606). The Java runtime system on server 312 adds the resulting hash value 404 and a pointer to the method 406 to mapping table 320 (step 608). When adding the hash value, the Java runtime system checks the mapping table to ensure that the hash value does not already exist in the mapping table, *i.e.*, no collisions have occurred with respect to the hash values. Although hash functions virtually guarantee that a hash value will uniquely identify a remote method, checking the table verifies that there are no collisions of hash values.

To illustrate an example of the steps used in Figure 6, suppose a remote object is created containing the following methods:

```
public interface Directory {  
    Address lookupAddress(String name);
```

```
Address lookupAddress(Person person)
}
```

The Java runtime system on server 312 creates a hash for each remote method. In this example, it creates two hashes:

Hash (lookupAddress, String), and
Hash (lookupAddress, Person).

Each hash value is unique and will be used to uniquely identify the remote method. Each hash value is added with a pointer to its corresponding method to mapping table 320, thus creating a method and hash value pairing 402 in mapping table 320. Server 312 can later access mapping table 320 using hash value 308 from client 302 to identify remote method 316 to be invoked.

The process of using hashes to identify remote methods on a remote server advantageously enables a client to uniquely identify the remote method without identifying an incorrect method. Additionally, the use of hashes avoids the need for long strings to identify remote methods, thereby allowing more efficient processing. The false identification of remote methods on servers commonly results from remote methods having string names common to more than one method, or the changing of numbering of methods without notifying clients using an old stub of the number changes. Methods and systems consistent with the present invention using hashes to identify remote methods on a remote server avoid these and related problems.

It will be appreciated by those skilled in this art that various modifications and variations can be made to the remote method identification strategy consistent with the present invention described herein without departing from the spirit and scope of the invention. Other embodiments of the invention will be apparent to those skilled in this art from consideration of the specification and practice of the invention disclosed herein. It is intended that the specification and examples be considered exemplary only, with a true scope and spirit of the invention being indicated by the following claims.

WHAT IS CLAIMED IS:

1. A method in a data processing system for invoking remote methods comprising the steps of:

providing an identifier uniquely identifying a remote method;
sending the identifier in response to an instruction to invoke the remote method; and
invoking the remote method based on the identifier.

2. The method of claim 1, wherein the invoking step includes the step of:
locating the remote method in a table indicating hash values and corresponding remote methods.

3. The method of claim 1, further including the step of:
returning a result of the invocation of the remote method.

4. The method of claim 1, wherein the providing step includes:
applying a hash function to an identifier and set of parameters for the remote method.

5. The method of claim 4, wherein applying step includes the step of:
applying an SHA-1 hash function.

6. The method of claim 1, wherein the invoking step includes the step of:
accessing a hash table containing hash values and references to remote methods.

7. A method in a data processing system for invoking remote methods comprising the steps of:

providing a hash value identifying a remote method;
sending the hash value in response to an instruction to invoke the remote method; and
invoking the remote method based on the hash value.

8. A method in a data processing system for identifying a remote method on a server by using a hash value comprising the steps of:

- providing a hash value uniquely identifying a remote method;
- sending the hash value in response to an instruction to invoke the remote method; and
- receiving a result of the invocation of the remote method.

9. The method of claim 8, wherein the providing step includes the step of:
applying a hash function to an identifier of the remote method and a corresponding parameter list to create the hash value.

10. The method of claim 9, wherein the applying step further includes the step of:
using a SHA-1 hash function.

11. A method in a data processing system for uniquely identifying methods in a distributed system comprised of a plurality of machines, the method comprising the steps, performed by one of the machines, of:

receiving a hash value identifying a method to be invoked;
identifying the method based on the hash value; and
performing the identified method.

5

12. The method of claim 11, wherein the performing step includes the step of:
invoking the identified method; and
returning a result of the invocation of the identified method.

13. The method of claim 11, wherein the identifying step includes the step of:
locating the method in a table using the hash value.

14. A system for identifying a method using a hash value, and in response, invoking the method, the system comprising:

a memory including:

a mapping table for mapping hash values to methods; and

5 a processor for:

receiving a request to invoke a selected method, the request including a hash value identifying the selected method;

accessing the mapping table to identify the selected method based on the hash value of the request; and

10 invoking the selected method.

15. A method for creating a mapping table for storing associations between hash values and remote methods comprises the steps of:

- creating a remote object;
- identifying a remote method implemented by the remote object;
- computing a hash value representing the remote method; and
- adding the hash value and the remote method to the mapping table.

16. The method of claim 15, wherein the identifying step includes the step of: identifying all of the remote methods implemented by the remote object.

17. The method of claim 15, wherein the computing step includes the step of: applying a hash function to the method name and parameter type list to compute the hash value.

18. A computer program product comprising a computer readable medium having computer readable code embodied therein for invoking remote methods by:

providing an identifier uniquely identifying a remote method;
sending the identifier in response to an instruction to invoke the remote method; and
invoking the remote method based on the identifier.

19. The product of claim 18, wherein the invoking step includes the step of:
locating the remote method in a table indicating hash values and corresponding remote methods.

20. The product of claim 18, further including the step of:
returning a result of the invocation of the remote method.

21. The product of claim 18, wherein the providing step includes:
applying a hash function to an identifier and set of parameters for the remote method.

22. A computer program product comprising a computer readable medium having computer readable code embodied therein for invoking remote methods by:

providing a hash value uniquely identifying a remote method;
sending the hash value in response to an instruction to invoke the remote method; and
receiving a result of the invocation of the remote method.

23. A computer program product comprising a computer readable medium having computer readable code embodied therein for invoking remote methods by:

receiving a hash value identifying a method to be invoked;
identifying the method based on the hash value; and
performing the identified method.

24. A computer program product comprising a computer readable medium having computer readable code embodied therein for invoking remote methods by:

creating a remote object;
identifying a remote method implemented by the remote object;
computing a hash value representing the remote method; and
adding the hash value and the remote method to the mapping table.

25. An apparatus in a data processing system for invoking remote methods comprising:

means for providing an identifier uniquely identifying a remote method;

5 means for sending the identifier in response to an instruction to invoke the remote method; and

means for invoking the remote method based on the identifier.





FIG. 1

2 / 6

**FIG. 2**




FIG. 3

4 / 6

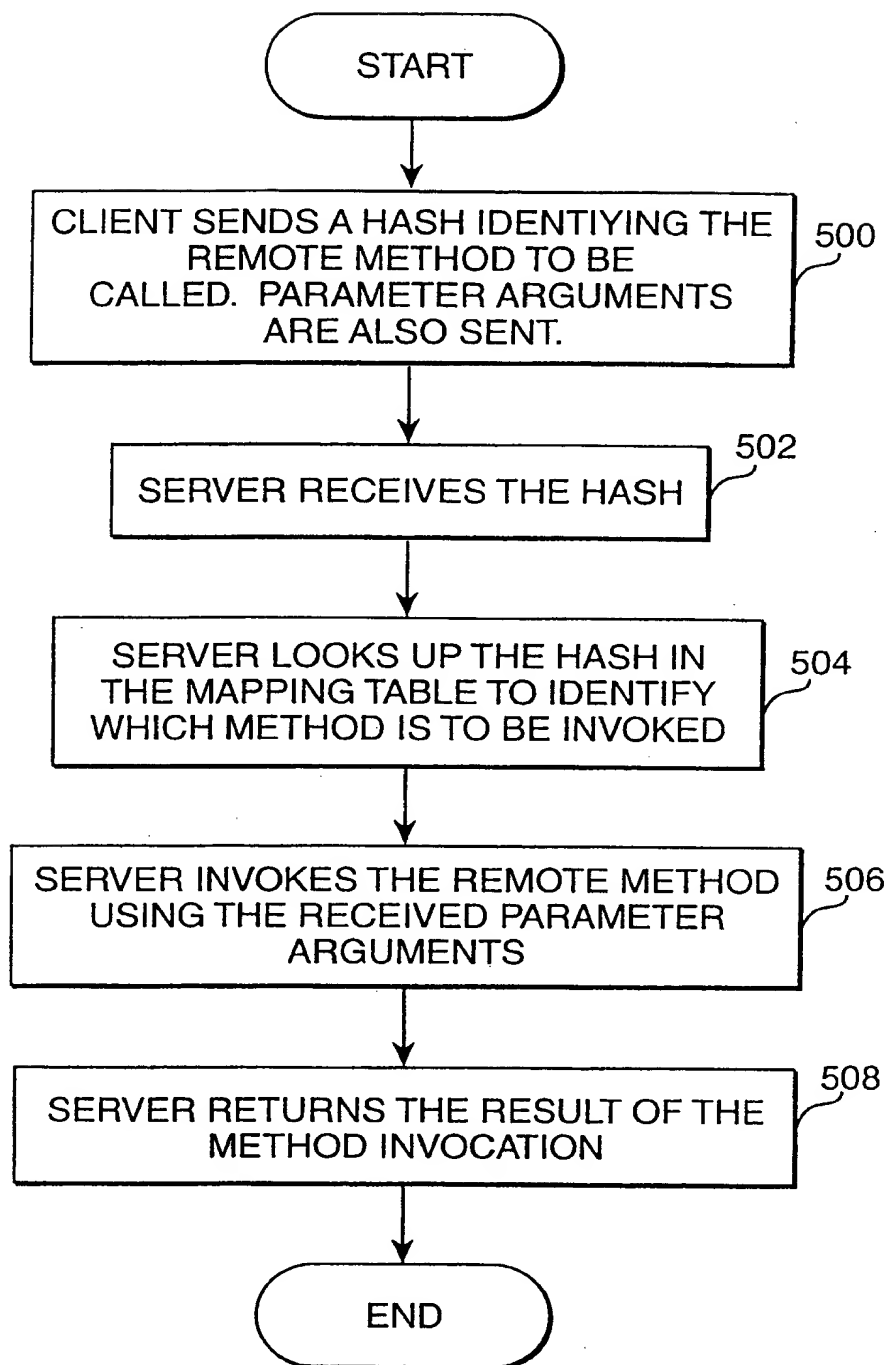
MAPPING TABLE FOR A SERVER CLASS

402 METHOD/ HASH PAIR	404 HASH VALUE	406 POINTER TO METHOD
	HASH VALUE	POINTER TO METHOD
	HASH VALUE	POINTER TO METHOD
	HASH VALUE	POINTER TO METHOD
	• • •	• • •

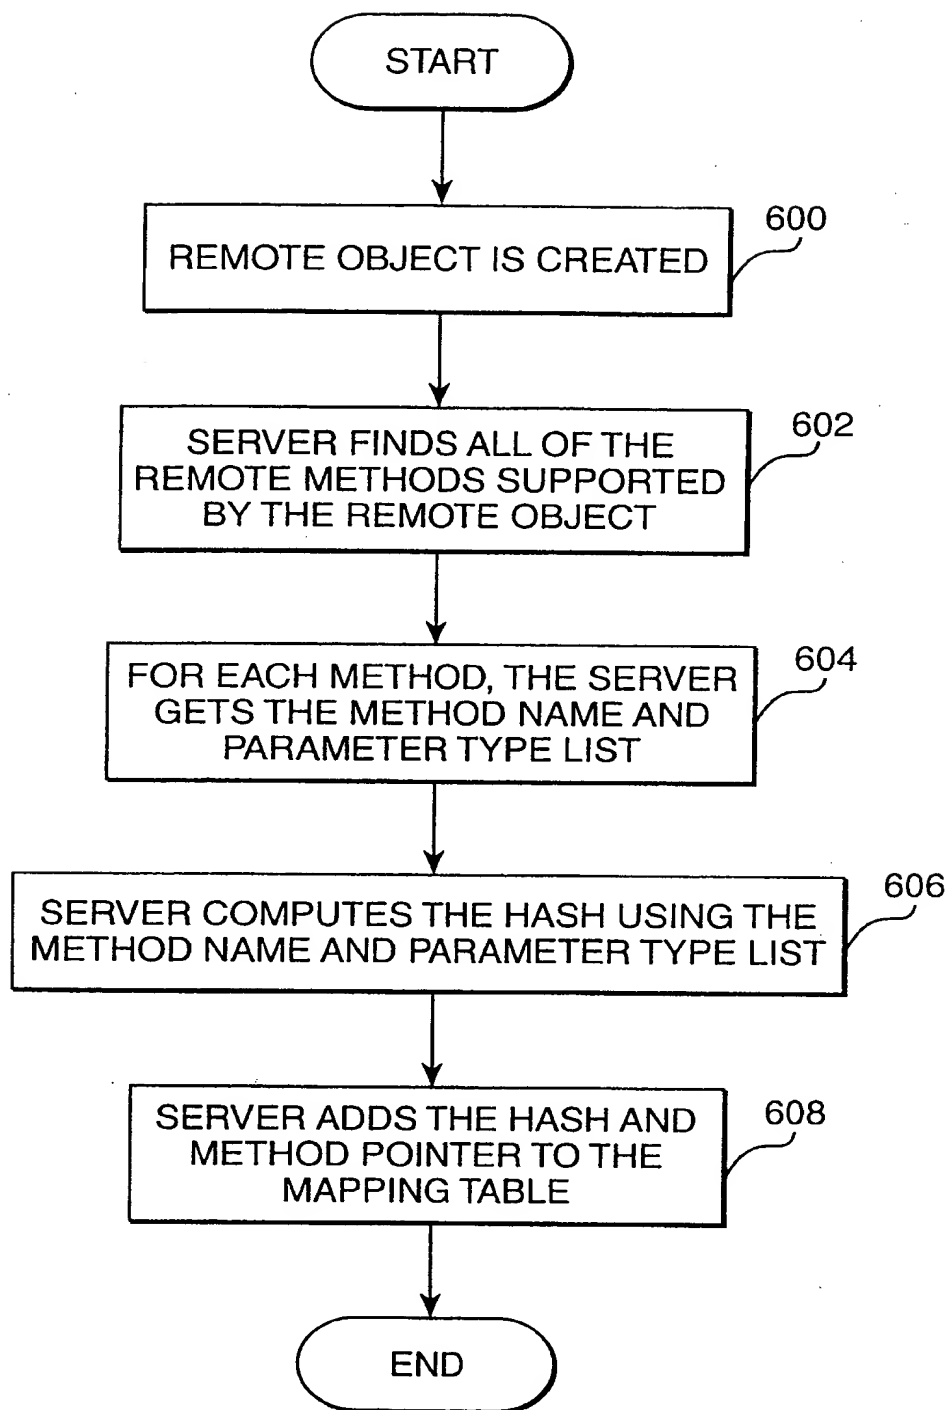
320

FIG. 4

5 / 6

**FIG. 5**

6 / 6

**FIG. 6**

SUBSTITUTE SHEET (RULE 26)